arXiv:1608.01426v1 [cs.CC] 4 Aug 2016

Solving Laplacian Systems in Logarithmic Space

François Le Gall
Graduate School of Informatics
Kyoto University
legall@i.u-kyoto.ac.jp

Abstract

We investigate the space complexity of solving linear systems of equations. While all known deterministic or randomized algorithms solving a square system of n linear equations in n variables require $\Omega(\log^2 n)$ space, Ta-Shma (STOC 2013) recently showed that on a quantum computer an approximate solution can be computed in logarithmic space, giving the first explicit computational task for which quantum computation seems to outperform classical computation with respect to space complexity. In this paper we show that for systems of linear equations in the Laplacian matrix of graphs, the same logarithmic space complexity can actually be achieved by a classical (i.e., non-quantum) algorithm. More precisely, given a system of linear equations $Lx = b$, where L is the (normalized) Laplacian matrix of a graph on n vertices and b is a unit-norm vector, our algorithm outputs a vector $\tilde{x}$ such that $\|\tilde{x} - x\| \leq 1/\mathrm{poly}(n)$ and uses only $O(\log n)$ space if the underlying graph has polynomially bounded weights. We also show how to estimate, again in logarithmic space, the smallest non-zero eigenvalue of L.

1 Introduction

Background. Ta-Shma showed a few years ago that several fundamental tasks in linear algebra, such as computing the eigenvalues of an $n \times n$ matrix A or finding a solution to a linear system of equations $Ax = b$, can be solved approximately with polynomial precision on a quantum computer using $O(\log n)$ space [16]. In comparaison, no $o(\log^2 n)$-space classical (i.e., non-quantum) algorithm is known for these problems: the best (in the space complexity setting) known classical algorithms that solve them exactly use $O(\log^2 n)$ space [1, 2, 5] and nothing better is known for only approximating the solutions. This breakthrough was the first example of concrete computational tasks for which a quantum algorithm outperforms the best known classical algorithms in the standard space complexity setting. The power of space-bounded quantum algorithms for matrix problems comes from their ability to represent an n-dimensional vector using $O(\log n)$ quantum bits of memory and consequently estimate the entries of A^k, given an $n \times n$ matrix A and an integer $k \leq \text{poly}(n)$, in logarithmic quantum space. In comparison, the best known algorithm for this task uses $O(\log^2 n)$ space [2].

Doron and Ta-Shma [6] recently showed how to "dequantize" the algorithm [16] for classes of matrices where A^k can be computed space efficiently, and in particular when A is the adjacency matrix of a graph, in which case entries of A^k can be estimated by a classical algorithm running k steps of the random walk corresponding to k. They focused on one task, computation of the eigenvalues, and succeeded in constructing a $O(\log n)$-space classical algorithm that computes the eigenvalues of such a matrix A, but only with constant precision (a much weaker accuracy that the polynomial precision obtained in the quantum case).

Our results. In this paper we investigate the classical (i.e., non-quantum) space complexity of matrix problems for matrices associated with graphs. More precisely, we focus on the Laplacian matrices of undirected weighted graphs (defined in Section 2).

The following definition will be useful to state our results concisely: We say that an undirected weighted graph G on n vertices has polynomially bounded weights if the weight of any edge is upper bounded by a polynomial of n and lower bounded by a polynomial of $1/n$. Unweighted graphs, where the weight of each edge is one, are examples of such graphs.

Our first result is stated in the following theorem.

Theorem 1.1. *Let G be an undirected weighted graph on n vertices with polynomially bounded weights. For any $\epsilon, \gamma \in (0,1)$, an ϵ-additive approximation of the solution of a Laplacian system corresponding to G can be computed with probability at least $1-\gamma$ in $O(\log(n/\epsilon) + \log\log(1/\gamma))$ space.*

Theorem 1.1 shows that if G has polynomially bounded weights, linear systems of equations involving the Laplacian matrix of G can be solved approximately with polynomial precision in $O(\log n)$ space, which gives a classical algorithm with the same space complexity and the same precision as Ta-Shma's quantum algorithm [16]. Note that Laplacian systems are a natural and very well-studied subclass of linear systems, and have a multitude of algorithmic applications (see [17] for a survey). While in the time complexity setting extremely fast algorithms for Laplacian systems have been obtained in the past decade [8, 11, 12, 13, 14, 15], we are not aware of any work on the space complexity of this problem.

Roughly speaking, Theorem 1.1 is proved by showing that solving a Laplacian system reduces to computing powers of the normalized adjacency matrix of the graph, and then computing

space-efficiently the entries of the powers by using random walks. We note that Chung and Simpson [4] considered a fairly similar approach based on random walks to design time-efficient algorithms solving Laplacian systems under boundary conditions. In the present paper the focus is different since we consider space complexity. Indeed, the main difficulty here is to show that the reduction to random walks can be implemented space efficiently.

The values of the eigenvalues of the Laplacian of a graph are related to many important properties of the graph — these relations are precisely the main subject of the field of spectral graph theory. The smallest non-zero eigenvalue is especially a fundamental quantity. If the graph is connected, then this eigenvalue is called the algebraic connectivity of the graph and controls for example expanding properties of the graph. Our second result, stated in the following theorem, shows that the smallest non-zero eigenvalue of the Laplacian matrix of an undirected weighted graph with polynomially bounded weights can be approximated with constant multiplicative precision in logarithmic space.

Theorem 1.2. *Let G be an undirected weighted graph on n vertices with polynomially bounded weights. For any $\delta, \gamma \in (0,1)$, a δ-multiplicative approximation of the smallest non-zero eigenvalue of the Laplacian matrix of G can be computed with probability at least $1-\gamma$ in $O(\frac{1}{\delta}\log n + \log\log(1/\gamma))$ space.*

This result is obtained by developing a space-efficient version of a well known technique, called the power method and used for estimating the largest eigenvalue of a matrix, and showing how to implementing it space-efficiently using random walks.

To our knowledge the only related prior work on the space complexity of approximating the algebraic connectivity of a graph is the recent work by Doron and Ta-Shma [6] mentioned above. Their techniques can be used to estimate the algebraic connectivity in logarithmic space with constant additive precision. Our techniques give constant multiplicative precision, which is a much stronger result (the smallest non-zero eigenvalue is always smaller than two; in many applications this eigenvalue is actually close to zero).

2 Preliminaries

General notations. In this paper, log always denote the natural logarithm. We use $\mathcal{P}_s$ to represent the Poisson distribution with parameter s (for any real number $s > 0$). This discrete probability distribution is defined as $\mathcal{P}_s(k) = \frac{e^{-s}s^s}{k!}$ for any integer $k \geq 0$.

For any real number $a < b$, the notation (a,b), respectively $(a,b]$, represents the set of real numbers x such that $a < x < b$, respectively $a < x \leq b$. Given two real numbers a, b and a positive real number δ, we say that a is a δ-additive approximation of b if $|a-b| \leq \delta$, and say that a is a δ-multiplicative approximation of b if $|a-b| \leq \delta|a-b|$.

We will generally work in the vector space $\mathbb{R}^n$ for some positive integer n (representing the number of vertices of the graph considered). Given a vector $v \in \mathbb{R}^n$, we write $\|v\|$ its Euclidean norm, and v^t its transpose. Given two vectors $u, v \in \mathbb{R}^n$, we write $\langle u, v\rangle$ their inner product. Given any $s \in \{1, \ldots, n\}$, we use e_s to denote the vector in $\mathbb{R}^n$ with s-th coordinate 1 and all other coordinates zero.

The Laplacian and its eigenvalues. We now define the Laplacian of a graph and introduce more specific notations used through this paper. It will be more convenient for us to work with the normalized version of the Laplacian of a graph, due to its natural connections with random walks. We refer to [3] for details on the normalized Laplacian and all the notions below.

Let $G = (V, E)$ be an undirected weighted graph (possibly with loops) on n vertices, i.e., a graph with a weight function $w\colon V \times V \to \mathbb{R}$ satisfying $w(i,j) = w(j,i)$ for all pairs of vertices $(i,j) \in V \times V$, $w(i,j) = 0$ if $(i,j) \notin E$ and $w(i,j) > 0$ if $(i,j) \in E$. The degree of a vertex $i \in V$, denoted d_i, is defined as $d_i = \sum_{j \in V} w(i,j)$. The degree matrix of the graph G is the $n \times n$ diagonal matrix $D_G = \mathsf{diag}(d_1, \ldots, d_n)$. The volume of G is defined as $\mathsf{vol}(G) = \sum_{\ell=1}^{n} d_\ell$. Finally, if the graph has no isolated vertex, i.e., the degree of each vertex is positive, let us define the quantity

$$\mathfrak{d} = \frac{\max_{i \in \{1,\ldots,n\}} d_i}{\min_{i \in \{1,\ldots,n\}} d_i}. \tag{1}$$

The (normalized) Laplacian of the graph G is the $n \times n$ symmetric matrix L_G such that

$$L_G[i,j] = \begin{cases} 1 - \frac{w(i,i)}{d_i} & \text{if } i = j \text{ and } d_i \neq 0, \\ -\frac{w(i,j)}{\sqrt{d_i d_j}} & \text{if } (i,j) \in E, \\ 0 & \text{otherwise,} \end{cases}$$

for all $i, j \in \{1, \ldots, n\}$. Let $\lambda_1 \leq \lambda_2 \cdots \leq \lambda_n$ denote its eigenvalues. It is known that $\lambda_1 \geq 0$ and $\lambda_n \leq 2$. Moreover, $\lambda_2 \neq 0$ if and only if G is a connected graph. We have the general upper bound $\lambda_2 \leq n/(n-1)$. When G is connected, we obtain the lower bound

$$\lambda_2 \geq \frac{1}{\mathsf{diam}(G)\mathsf{vol}(G)}, \tag{2}$$

where $\mathsf{diam}(G)$ denotes the diameter of G (the maximal distance between two vertices in G). Let $\{u_1, \ldots, u_n\}$ be an orthonormal basis for L_G, where u_i is an eigenvector associated with eigenvalue λ_i. We can thus write $L_G = \sum_{i=1}^{n} \lambda_i u_i u_i^t$. It is easy to check that $u_1 = \frac{1}{\sqrt{\mathsf{vol}(G)}} \left(\sqrt{d_1}, \ldots, \sqrt{d_n}\right)^t$.

Image of L_G and its pseudo inverse. We define $\mathsf{Im}(L_G)$ as the image of L_G, that is, the linear span of the vectors $u_c, u_{c+1}, \ldots, u_n$, where c is the smallest integer such that $\lambda_c \neq 0$. The pseudo inverse (also called the Moore-Penrose pseudo inverse) of L_G, denoted $L_G^\dagger$, is the matrix

$$L_G^\dagger = \sum_{i=c}^{n} \lambda_i^{-1} u_i u_i^t,$$

which corresponds to the inverse of L_G when restricted to the subspace $\mathsf{Im}(L_G)$. Note that if the graph G is connected then $c = 2$. In that case $\mathsf{Im}(L_G)$ is simply the subspace of $\mathbb{R}^n$ consisting of all vectors orthogonal to u_1.

The transition matrix and random walks. Another useful matrix is the transition matrix of G, which represents one step of a random walk on the graph G. Assume that the graph has no isolated vertex. Then the transition matrix of G is the $n \times n$ matrix P_G defined as

$$P_G[i,j] = \begin{cases} w(i,j)/d_i & \text{if } (i,j) \in E, \\ 0 & \text{otherwise} \end{cases}$$

for all $i, j \in \{1, \ldots, n\}$. Note that $L_G = I - D_G^{1/2} P_G D_G^{-1/2}$.

Laplacian systems. Given a Laplacian L_G of an undirected weighted graph G and a vector $b \in \mathsf{Im}(L_G)$, the corresponding Laplacian system is the equation $L_G x = b$. The goal is to solve this equation, i.e., compute $L_G^\dagger b$ or to compute an (additive or approximative) approximation of

it. A standard assumption is that G is connected since one can always reduce to this case dealing individually with each connected component (note that this reduction is space-efficient since one can identify the connected components of a graph in logarithmic space). For this reason, in Sections 3 and 4 of this paper we will assume that G is connected. For convenience, and without loss of generality, we will also assume that b is a unit-norm vector.

Remark on numerical precision. The algorithms we present in this paper perform arithmetic operations on real numbers. When working with real numbers numerical precision is always a delicate issue, and especially in a space-bounded setting. To illustrate this point, let us consider as a simple example the multiplication of an integer a by the irrational numbers $\sqrt{2}$. One of the most rigorous ways to proceed is to work in the bit complexity model, and compute only a fixed number of bits of $a\sqrt{2}$, which necessarily introduce a (small) approximation error. Analyzing rigorously a complicated algorithm in this model is typically extremely tedious, since the approximation errors of essentially all the arithmetic steps of the algorithm have to been considered. In this paper we (implicitly) use a slightly more abstract, but still fairly standard, model where basic arithmetic operations involving "reasonable" numbers can be implemented exactly. More precisely, we assume that the standard arithmetic operations (addition, substraction, multiplication, division) on $O(\log m)$-bit reals numbers (i.e., real numbers of absolute value between m^{-c} and $m^{c'}$ for some constant c and c') can be done exactly in $O(\log m)$ space. This model enables us to focus on the most interesting and important algorithmic aspects of our approach, without having to deal with a multitude of minor technical details. Note that all the statements of the technical results of Sections 3 and 4 hold with only minor modifications in the bit complexity model as well. The statements of Theorems 1.1 and 1.2 in the introduction, described in the more convenient setting of graphs with polynomially bounded weighted, hold without any modification even in the bit complexity model.

3 Space-efficient Laplacian Solver

In this section we describe how to compute in logarithmic space an approximate solution of the Laplacian system $L_G x = b$, and prove Theorem 1.1. We assume that G is connected.

The theoretical foundation of our algorithm comes from the formula of the following theorem, which shows that the solution $L_G^\dagger b$ to the Laplacian system can be approximated using powers of the transition matrix P_G of the graph.

Theorem 3.1. *Let ϵ be any positive real number and $\lambda \in (0, 2]$ be a lower bound on all the non-zero eigenvalues of L_G (i.e., a lower bound on λ_2). For any positive integers $T \geq \frac{\log(6/(\epsilon\lambda))}{\lambda}$, $N \geq \frac{6T}{\epsilon}$ and $K \geq \max(6T, \log(6T/\epsilon))$, the inequality*

$$\left\| \left(L_G^\dagger - \frac{T}{N} \sum_{j=1}^{N} \sum_{k=0}^{K-1} \mathcal{P}_{jT/N}(k) D_G^{1/2} P_G^k D_G^{-1/2} \right) b \right\| \leq \frac{\epsilon \|b\|}{2}$$

holds for any vector $b \in \mathsf{Im}(L_G)$.

Proof. Observe that $D_G^{1/2} P_G^k D_G^{-1/2} = \sum_{i=1}^{n} (1-\lambda_i)^k u_i^t u_i$ for any integer $k \geq 0$. We show below that the inequality

$$\left| \frac{1}{\lambda_i} - \frac{T}{N} \sum_{j=1}^{N} \sum_{k=0}^{K-1} \mathcal{P}_{jT/N}(k) (1-\lambda_i)^k \right| \leq \epsilon/2,$$

holds for all $i \in \{2, \ldots, n\}$, which will imply the statement in the theorem.

For any $i \in \{2, \ldots, n\}$, first observe that $\frac{1}{\lambda_i} = \int_0^\infty e^{-t\lambda_i} dt$. For any $T \geq 0$, we thus have

$$\left| \frac{1}{\lambda_i} - \int_0^T e^{-t\lambda_i} dt \right| \leq \int_T^\infty e^{-t\lambda_i} dt = \frac{e^{-T\lambda_i}}{\lambda_i}. \tag{3}$$

Let us approximate the integral by a right Riemann sum: for any integer $N \geq 1$,

$$\left| \int_0^T e^{-t\lambda_i} dt - \frac{T}{N} \sum_{j=1}^{N} e^{-j\frac{T}{N}\lambda_i} \right| \leq \frac{T}{N} \times (1 - e^{-T\lambda_i}) \leq \frac{T}{N}. \tag{4}$$

For any $s \geq 0$ we have

$$e^{-s\lambda_i} = e^{-s} e^{(1-\lambda_i)s} = e^{-s} \sum_{k=0}^{\infty} \frac{((1-\lambda_i)s)^k}{k!} = \sum_{k=0}^{\infty} \mathcal{P}_s(k)(1-\lambda_i)^k.$$

For any integer $K \geq 1$ we thus have

$$\left| e^{-s\lambda_i} - \sum_{k=0}^{K-1} \mathcal{P}_s(k)(1-\lambda_i)^k \right| \leq \sum_{k=K}^{\infty} \mathcal{P}_s(k) \leq \left(1 - \frac{s}{K+1}\right)^{-1} \mathcal{P}_s(K) \leq \left(1 - \frac{s}{K+1}\right)^{-1} \frac{s^K e^{-s}}{\sqrt{2\pi K} K^K e^{-K}}.$$

The first inequality uses $\lambda_i \in (0, 2]$. The second inequality uses a well-known upper bound on the tail probability of the Poisson distribution [7]. The third inequality uses a standard lower bound on $K!$ from Stirling approximation. For any positive integer $K \geq e^2 s$, where $e = 2.718\ldots$ is Euler's number, and in particular for any $K \geq 6s$, we get

$$\left| e^{-s\lambda_i} - \sum_{k=0}^{K-1} \mathcal{P}_s(k)(1-\lambda_i)^k \right| \leq \left(1 - \frac{1}{e^2}\right)^{-1} \frac{e^{-s}}{\sqrt{2\pi K} e^K} \leq \frac{1}{e^K}. \tag{5}$$

Combining Inequalities (3), (4) and (5), for any positive integer $K \geq 6T$ we obtain:

$$\left| \frac{1}{\lambda_i} - \frac{T}{N} \sum_{j=1}^{N} \sum_{k=0}^{K-1} \mathcal{P}_{jT/N}(k)(1-\lambda_i)^k \right| \leq \frac{e^{-T\lambda_i}}{\lambda_i} + \frac{T}{N} + \frac{T}{e^K}.$$

Taking values of T, N and K as in the statement of the theorem guarantees that the right side of the above inequality is at most $\epsilon/2$. □

We now present a lemma that shows how to approximate space-efficiently the Poisson distribution appearing in the formula of Theorem 3.1. This is done by using the well-known property that the Poisson distribution can be expressed as the limit distribution of binomial random variables, and showing that the convergence is fast enough.

Lemma 3.2. *For any $\delta \in (0, 1)$ and any $\zeta > 0$, there exists a $O(\log(ks/\delta) + \log\log(1/\zeta))$-space algorithm that outputs a δ-additive approximation of $\mathcal{P}_s(k)$ with probability at least $1 - \zeta$.*

Proof. Let n be any integer such that $n \geq 2(k^2+s^2)/\delta$. Let X_n be the binomial random variable with parameters n and s/n, i.e., X_n is the number of successes in n repeated trials of a binomial experiment with success probability s/n. Standard computations (see, e.g., page 99 of [10]) show that

$$\mathcal{P}_s(k)(1-k/n)^k\left(1-\frac{s^2}{n}\right) \leq \Pr[X_n = k] \leq \mathcal{P}_s(k)\left(\frac{1}{1-sk/n}\right).$$

Using the inequality $1/(1-x) \leq 1+2x$ valid for any $x \in (0,1/2)$ and the inequality $(1-x)^k \geq 1-kx$ valid for any $k \geq 0$, we get

$$\mathcal{P}_s(k)\left(1-\frac{k^2+s^2}{n}\right) \leq \Pr[X_n = k] \leq \mathcal{P}_s(k)\left(1+\frac{2sk}{n}\right).$$

With our value of n, and since $\mathcal{P}_s(k) \leq 1$, we obtain

$$|\mathcal{P}_s(k) - \Pr[X_n = k]| \leq \delta/2. \tag{6}$$

Our algorithm for estimating $\mathcal{P}_s(k)$ is as follows. We will use the value $n = \lceil 2(k^2+s^2)/\delta \rceil$. The algorithm creates a counter C initialized to zero, and repeat $m = \left\lceil \frac{2\log(2/\zeta)}{\delta^2} \right\rceil$ times the following: perform n trials of a binomial experiment with success probability s/n and increment C by one if exactly k among these n trials succeeded. The algorithm finally outputs C/m.

The expected value of the output of this algorithm is precisely the quantity $\Pr[X_n = k]$ considered above. From Chernoff bound, the output of the algorithm is thus a $\delta/2$-additive approximation of this quantity with probability at least $1 - 2e^{-m\delta^2/2} \geq 1-\zeta$. Combining this with Inequality (6) and the triangular inequality, we conclude that the output of the algorithm is a δ-additive approximation of $\mathcal{P}_s(k)$ with the same probability.

Finally, observe that the space complexity of the algorithm is linear in $\log(1/\delta)$, $\log k$, $\log s$ and $\log\log(1/\zeta)$. □

We are now ready to present our algorithm for estimating $L_G^\dagger b$. The algorithm, denoted Algorithm $\mathcal{A}$, is described in Figure 1 and analyzed in the following theorem.

Theorem 3.3. *Let G be an undirected weighted connected graph and $\lambda \in (0,2]$ be a lower bound on the second smallest zero eigenvalue λ_2 of G. Algorithm $\mathcal{A}$ outputs an ϵ-additive approximation of the i-th entry of $L_G^\dagger b$ with probability at least $1-\gamma$, and uses $O(\log\left(\frac{n\mathfrak{d}}{\epsilon\lambda}\right) + \log\log(1/\gamma))$ space.*

Proof. Remember that for any $s \in \{1,\ldots,n\}$, e_s denotes the $1 \times n$ vector with s-th coordinate 1 and all other coordinates zero. Note that the i-th coordinate of $L_G^\dagger b$, which is the quantity we want to estimate, is $\left\langle e_i, b^t L_G^\dagger \right\rangle$ since $L_G^\dagger$ is symmetric. From Theorem 3.1 and the triangular inequality, we know that any $\frac{\epsilon}{2}$-additive approximation of

$$\frac{T}{N}\sum_{j=1}^{N}\sum_{k=0}^{K-1}\mathcal{P}_{jT/N}(k)\left\langle e_i, b^t D_G^{1/2} P_G^k D_G^{-1/2}\right\rangle = \frac{T}{N}\sum_{j=1}^{N}\sum_{k=0}^{K-1}\sum_{\ell=1}^{n}\mathcal{P}_{jT/N}(k) b_\ell \sqrt{\frac{d_\ell}{d_i}}\left\langle e_i, e_\ell P_G^k\right\rangle$$

is an ϵ-additive approximation of $\left\langle e_i, b^t L_G^\dagger\right\rangle$. We show below that Algorithm $\mathcal{A}$ precisely outputs an $\epsilon/2$-additive approximation of this quantity.

Input: an undirected, weighted and connected graph G on n vertices,
an integer $i \in \{1, \ldots, n\}$, a unit vector $b \in \mathsf{Im}(L_G)$
a precision parameter $\epsilon \in (0, 1]$, an error parameter $\gamma \in (0, 1]$

1 $T \leftarrow \left\lceil \frac{\log(6/(\epsilon\lambda))}{\lambda} \right\rceil$; $N \leftarrow \lceil \frac{6T}{\epsilon} \rceil$; $K \leftarrow \lceil \max(6T, \log(\frac{6T}{\epsilon})) \rceil$;
2 $\delta \leftarrow \left(6TK\sqrt{\sum_{\ell=1}^{n} \frac{d_\ell}{d_i}}\right)^{-1}$; $\zeta \leftarrow \frac{\gamma}{NK(1+n)}$; $r \leftarrow \left\lceil \frac{\log(2/\zeta)}{2\delta^2} \right\rceil$;
3 $R \leftarrow 0$;
4 **for** j from 1 to N
5 **for** k from 0 to $K-1$
6 Compute an approximation a of $\mathcal{P}_{jT/N}(k)$ using Lemma 3.2 with δ and ζ;
7 **for** ℓ from 1 to n
8 $S \leftarrow 0$;
9 **repeat** r times:
10 Run the walk P_G starting on vertex ℓ for k steps;
11 If the walk ends on vertex i then $S \leftarrow S + 1$;
12 $R \leftarrow R + \frac{ab_\ell S}{r} \times \sqrt{\frac{d_\ell}{d_i}}$;
13 Output RT/N.

Figure 1: Algorithm $\mathcal{A}$ computing an ϵ-additive approximation of the i-th entry of $L_G^\dagger b$ with probability at least $1 - \gamma$.

Note that the probability of a walk P_G starting on vertex ℓ reaches vertex i after exactly k steps is $\langle e_i, e_\ell P_G^k \rangle$, the i-th coordinate of $e_\ell P_G^k$. At the end of Steps 9-11 we thus have

$$\Pr\left[|S/r - \langle e_i, e_\ell P_G^k \rangle| \leq \delta\right] \geq 1 - 2e^{-2r\delta^2} \geq 1 - \zeta, \tag{7}$$

from Chernoff bound. Lemma 3.2 also shows that a is a δ-additive approximation of $\mathcal{P}_{jT/N}(k)$ with probability at least $1 - \zeta$. Let us continue our analysis under the assumption that all these approximations are correct (we discuss the overall success probability at the end of the proof). At Step 12 we thus have

$$\left| \frac{aS}{r} - \mathcal{P}_{jT/N}(k) \left\langle e_i, e_\ell P_G^k \right\rangle \right| \leq \left(\delta + \mathcal{P}_{jT/N}(k) + \left\langle e_i, e_\ell P_G^k \right\rangle \right) \delta \leq 3\delta.$$

The output of the algorithm at Step 13 then satisfies

$$\left| \frac{RT}{N} - \frac{T}{N} \sum_{j=1}^{N} \sum_{k=0}^{K-1} \sum_{\ell=1}^{n} \mathcal{P}_{jT/N}(k) b_\ell \sqrt{\frac{d_\ell}{d_i}} \left\langle e_i, e_\ell P_G^k \right\rangle \right| \leq \frac{T}{N} \sum_{j=1}^{N} \sum_{k=0}^{K-1} \sum_{\ell=1}^{n} 3\delta |b_\ell| \sqrt{\frac{d_\ell}{d_i}} = TK \sum_{\ell=1}^{n} 3\delta |b_\ell| \sqrt{\frac{d_\ell}{d_i}}$$

$$\leq 3\delta TK \sqrt{\sum_{\ell=1}^{n} \frac{d_\ell}{d_i}} \leq \epsilon/2.$$

The space complexity of the algorithm is $O(\log(n\mathfrak{d}/(\lambda\epsilon)) + \log\log(1/\gamma))$, from Lemma 3.2 and the observation that only registers of this size are needed to implement the algorithm. Finally,

let us discuss the success probability of this algorithm. Errors can only occur at Steps 6 or 10. From Inequality (7) and Lemma 3.2, and using the union bound, we know that the overall success probability is at least $1 - \zeta NK(1+n) = 1 - \gamma$. □

Theorem 3.3 implies Theorem 1.1 by observing that $\mathfrak{d}/\lambda$ can be upper bounded by a polynomial in n when the weights are polynomially bounded, as shown in Equations (1) and (2) of Section 2 (if G is not connected we can simply apply Theorem 4.4 on each connected component).

4 Space-efficient Approximation of the Spectral Gap

Let us consider the following matrix:

$$M_G = \frac{1}{2}\left(I + D_G^{1/2} P_G D_G^{-1/2}\right).$$

Note that M_G is a symmetric matrix. Its eigenvalues are $0 \leq 1 - \lambda_n/2 < \cdots < 1 - \lambda_2/2 < 1 - \lambda_1/2 = 1$. The eigenvectors of M_G are the same as the eigenvectors of L_G. In particular, the eigenvector of M_G corresponding to the eigenvalue 1 is u_1, and the eigenvector of M_G corresponding to the eigenvalue $1 - \lambda_2/2$ is u_2. A well-known approach for approximating the largest eigenvalue of a matrix is the power method (see, e.g., [17]). Our idea is to apply this method on M_G restricted to $\mathsf{Im}(L_G)$, and compute the ratio $\|M_G^{k+1}v\|/\|M_G^k v\|$ on a random vector $v \in \mathsf{Im}(L_G)$ — it is easy to show that with high probability this ratio is close to $1 - \lambda_2$ for large enough k. In this section we will develop a space-efficient version of this approach, and prove Theorem 1.2

When using the power method to estimate $1 - \lambda_2/2$, we need to use vectors orthogonal to u_1 that have a non-negliglible "overlap" with the eigenvector u_2. We say that a vector $v \in \mathbb{R}^n$ is good if the following three conditions are satisfied: $\|v\| = 1$, $v \in \mathsf{Im}(L_G)$, and $|\langle v, u_2\rangle| \geq \frac{1}{\sqrt{2n\mathfrak{d}}}$. While a random unit-norm vector in $\mathsf{Im}(L_G)$ is a good vector with high probability, several technical difficulties arise when considering space-efficient vector sampling. Instead of using such probabilistic arguments, we introduce below a set Σ of vectors that necessarily contains at least one good vector.

Let $\Sigma \subset \mathbb{R}^n$ be the set containing the $n(n-1)/2$ vectors defined as follows. Each of these vectors corresponds to choosing two distinct indexes $i, j \in \{1, \ldots, n\}$ and taking the n-dimensional vector with i-th coordinate $-\frac{1}{\sqrt{1+d_i/d_j}}$, j-th coordinate $\frac{1}{\sqrt{1+d_j/d_i}}$, and all other coordinates being zero. The following easy lemma shows that Σ indeed contains at least one good vector.

Lemma 4.1. *There exists a good vector in* Σ.

Proof. Any vector in Σ is a unit-norm vector orthogonal to u_1. We show below that there exists a vector in Σ that also satisfies the third condition of the definition of good vectors.

Let us write $u_2 = (x_1, \ldots, x_n)$. Let $S^+ \subseteq \{1, \ldots, n\}$ be the set of indices ℓ such that $x_\ell \geq 0$, and $S^- \subseteq \{1, \ldots, n\}$ be the set of indices ℓ such that $x_\ell < 0$. Since u_2 is orthogonal to u_1, we have $\sum_{\ell \in S^+} \sqrt{d_\ell}|x_\ell| = \sum_{\ell \in S^-} \sqrt{d_\ell}|x_\ell|$. Since u_2 is a unit vector we have

$$1 = \sum_{\ell \in S^+} x_\ell^2 + \sum_{\ell \in S^-} x_\ell^2 \leq \sum_{\ell \in S^+} |x_\ell| + \sum_{\ell \in S^-} |x_\ell|.$$

We conclude that

$$\sum_{\ell \in S^+} \sqrt{d_\ell}|x_\ell| \geq \frac{\min_{\ell \in \{1,\ldots,n\}} \sqrt{d_\ell}}{2} \text{ and } \sum_{\ell \in S^-} \sqrt{d_\ell}|x_\ell| \geq \frac{\min_{\ell \in \{1,\ldots,n\}} \sqrt{d_\ell}}{2}$$

which implies that there exist $i \in S^-$ and $j \in S^+$ such that $|x_i|$ and $|x_j|$ are at least $\frac{1}{2n\sqrt{\mathfrak{d}}}$. Let v be the vector in Σ with i-th coordinate $-\frac{1}{\sqrt{1+d_i/d_j}}$ and j-th coordinate $\frac{1}{\sqrt{1+d_j/d_i}}$. The inner product of v and u_2 is thus at least

$$\frac{1}{2n\sqrt{\mathfrak{d}}}\left(\frac{1}{\sqrt{1+d_i/d_j}}+\frac{1}{\sqrt{1+d_j/d_i}}\right) = \frac{1}{2n\sqrt{\mathfrak{d}}}\left(\frac{d_j}{\sqrt{d_j+d_i}}+\frac{d_i}{\sqrt{d_i+d_j}}\right) \geq \frac{1}{\sqrt{2}n\mathfrak{d}},$$

as claimed. □

The following proposition is our version of the power method.

Proposition 4.1. *Let δ be any real number such that $0 < \delta \leq 1$, and ζ be any real number such that $0 < \zeta \leq \delta\lambda_2/12$.*

(i) For any integer $k \geq 0$, any non-zero vector $v \in \mathsf{Im}(L_G)$ and any ζ-multiplicative approximations C_1 and C_2 of $\|M_G^k v\|$ and $\|M_G^{k+1} v\|$, respectively, the inequality $(1-\delta)\lambda_2 \leq 2\,(1 - C_2/C_1)$ holds.

(ii) Let v be a good vector. For any integer $k \geq \frac{3\log(\sqrt{2}n\mathfrak{d})}{\delta\lambda_2} - 1$ and any ζ-multiplicative approximations C_1 and C_2 of $\|M_G^k v\|$ and $\|M_G^{k+1} v\|$, respectively, the inequality $2\,(1 - C_2/C_1) \leq (1+\delta)\lambda_2$ holds

Proof. Let us first prove part (i). We have

$$\frac{C_2}{C_1} \leq \frac{\left\|M_G^{k+1} v\right\|}{\left\|M_G^k v\right\|} \times \frac{1+\zeta}{1-\zeta} \leq (1-\lambda_2/2) \times \frac{1+\zeta}{1-\zeta} = \left(1+\frac{2\zeta}{1-\zeta}\right)(1-\lambda_2/2)$$

and thus

$$2\left(1-\frac{C_2}{C_1}\right) \geq \lambda_2 - \frac{4\zeta}{1-\zeta} + \frac{2\zeta\lambda_2}{1-\zeta} \geq \left(1-\frac{4\zeta}{\lambda_2(1-\zeta)}\right)\lambda_2 \geq \left(1-\frac{5\zeta}{\lambda_2}\right)\lambda_2 \geq (1-\delta)\,\lambda_2,$$

where the third inequality was obtained from $\zeta \leq 1/5$ (from the assumption $\zeta \leq \delta\lambda_2/12$).

Let us now prove part (ii). For any unit vector v, Hölder's inequality gives

$$\left\|M_G^k v\right\|^2 \leq \left(\sum_i v_i^2 \lambda_i^{2k+2}\right)^{k/(k+1)} \left(\sum_i v_i^2\right)^{1/(k+1)} = \left\|M_G^{k+1} v\right\|^{2k/(k+1)}.$$

If v is good then

$$\frac{\|M_G^{k+1} v\|}{\|M_G^k v\|} \geq \frac{\|M_G^{k+1} v\|}{\|M_G^{k+1} v\|^{k/(k+1)}} = \|M_G^{k+1} v\|^{1/(k+1)} \geq \frac{(1-\lambda_2/2)}{(\sqrt{2}n\mathfrak{d})^{1/(k+1)}}.$$

Taking $k+1 \geq \frac{3\log(\sqrt{2}n\mathfrak{d})}{\delta\lambda_2}$ gives $\frac{1}{(\sqrt{2}n\mathfrak{d})^{1/(k+1)}} \geq e^{-\delta\lambda_2/3} \geq 1-\delta\lambda_2/3$. Let C_1 and C_2 be as in the statement of the proposition. We have

$$\begin{aligned}\frac{C_2}{C_1} \geq \frac{\|M_G^{k+1}v\|}{\|M_G^k v\|} \times \frac{1-\zeta}{1+\zeta} &\geq (1-\lambda_2/2)(1-\delta\lambda_2/3) \times (1-2\zeta) \\ &\geq 1 - 2\zeta - \frac{\delta\lambda_2}{3} + \frac{2\zeta\delta\lambda_2}{3} - \frac{\lambda_2}{2} + \lambda_2\zeta + \frac{\delta\lambda_2^2}{3}\left(\frac{1}{2}-\zeta\right) \\ &\geq 1 - 2\zeta - \frac{\delta\lambda_2}{3} - \frac{\lambda_2}{2},\end{aligned}$$

where the last inequality uses $\zeta \leq 1/2$ (which is guaranteed from the assumption $\zeta \leq \delta\lambda_2/12$). We thus obtain $2\left(1 - C_2/C_1\right) \leq \left(1 + 2\delta/3 + 4\zeta/\lambda_2\right)\lambda_2 \leq (1+\delta)\lambda_2$, as claimed. □

Proposition 4.1 requires good multiplicative approximations of $\|M_G^k v\|$ and $\|M_G^{k+1}v\|$ to approximate λ_2. Using random walks, we nevertheless will only be able to obtain additive approximations. To convert additive approximations into a good multiplicative approximations, we will need lower bounds on these two quantities. We will also need upper bounds in order to control the running time (and the space complexity) of our algorithm. The following lemma shows the bounds we will use.

Lemma 4.2. *Let τ be any real number such that $\tau \in (0,1]$. For any unit vector $v \in \mathsf{Im}(L_G)$, $\|M_G^k v\| < \tau$ for all integers $k > \frac{2\log(1/\tau)}{\lambda_2}$. Additionally, if v is good and $n \geq 4$ then $\|M_G^{k+1}v\| \geq 2\tau$ for all integers*

$$k \leq \frac{\log(1/\tau) - \log(2\sqrt{2}n\mathfrak{d})}{2\lambda_2} - 1. \tag{8}$$

Proof. For any unit vector $v \in \mathsf{Im}(L_G)$ we have

$$\left\|M_G^k v\right\| \leq (1-\lambda_2/2)^k \leq e^{-k\lambda_2/2},$$

which is upper bounded by τ for $k > \frac{2\log(1/\tau)}{\lambda_2}$. If v is good we further have

$$\left\|M_G^{k+1}v\right\| \geq \frac{(1-\lambda_2/2)^{k+1}}{\sqrt{2}n\mathfrak{d}} \geq \frac{1}{\sqrt{2}n\mathfrak{d}}\left(e^{-1} - \frac{\lambda_2}{4}\right)^{(k+1)\lambda_2/2},$$

where we used the formula $e^{-1} - \frac{1}{2a} \leq (1-1/a)^a$ valid for any $a \geq 1$ (see, e.g., [9]) with $a = 2/\lambda_2$. Note that

$$e^{-1} - \lambda_2/4 \geq e^{-1} - \frac{n}{4(n-1)} > \frac{3}{100}$$

for $n \geq 4$. We get $\left\|M_G^{k+1}v\right\| \geq 2\tau$ whenever

$$k+1 \leq \frac{2}{\lambda_2} \times \frac{\log(2\sqrt{2}\tau n\mathfrak{d})}{\log(3/100)} = \frac{2}{\log(100/3)} \times \frac{\log(1/\tau) - \log(2\sqrt{2}n\mathfrak{d})}{\lambda_2}.$$

Finally, note that $2/\log(100/3) > 1/2$. □

In order to use the the theory developed above, we need to be able to estimate $\left\|M_G^k v\right\|$ space-efficiently. This can be done using an approach based on quantum walks, similarly to what we did in Section 3. The description the procedure based on this idea, and its analysis, are given in the appendix. This procedure, denoted ESTIMATE-NORM(G,k,v,ϵ,γ), computes an ϵ-additive approximation of $\left\|M_G^k v\right\|$ with probability at least $1-\gamma$, for any $k \geq 0$, any $v \in \Sigma$ and any $\epsilon, \gamma \in (0,1]$. We state the main result of the appendix as the following theorem.

Theorem 4.3. *Procedure* ESTIMATE-NORM(G,k,v,ϵ,γ) *uses* $O(\log(nk\mathfrak{d}/\epsilon) + \log\log(1/\gamma))$ *space and outputs an* ϵ*-additive approximation of* $\left\|M_G^k v\right\|$ *with probability at least* $1-\gamma$.

Our algorithm for estimating λ_2, denoted Algorithm $\mathcal{B}$, is given in Figure 2 and analyzed in Theorem 4.4 below.

Input: an undirected, weighted and connected graph G on n vertices, where $n \geq 4$,
a precision parameter $\delta \in (0,1]$, an error parameter $\gamma \in (0,1]$

1 $\tau \leftarrow \frac{1}{2(\sqrt{2}n\mathfrak{d})^{1+8/\delta}}$; $\epsilon \leftarrow \frac{\delta\lambda\tau}{12}$; $\zeta \leftarrow \frac{4\gamma}{n(n-1)} \times \left(1 + \frac{\log(1/\tau)}{\lambda_2}\right)^{-1}$;
2 $R_{\max} \leftarrow 0$;
3 **for** all $v \in \Sigma$
4 $\quad k \leftarrow 1$;
5 $\quad C_1 \leftarrow$ ESTIMATE-NORM$(G,1,v,\epsilon,\zeta)$; # C_1 will store an approximation of $\left\|M_G^k v\right\|$
6 $\quad C_2 \leftarrow$ ESTIMATE-NORM$(G,2,v,\epsilon,\zeta)$; # C_2 will store an approximation of $\left\|M_G^{k+1} v\right\|$
7 $\quad$ **repeat** until $C_2 < 3\tau/2$
8 $\quad\quad$ **if** $C_2/C_1 > R_{\max}$
9 $\quad\quad\quad R_{\max} \leftarrow C_2/C_1$;
10 $\quad\quad k \leftarrow k+1$;
11 $\quad\quad C_1 \leftarrow C_2$;
12 $\quad\quad C_2 \leftarrow$ ESTIMATE-NORM$(G,k+1,v,\epsilon,\zeta)$;
13 Output $2(1-R_{\max})$.

Figure 2: Algorithm $\mathcal{B}$ computing an δ-multiplicative approximation of λ_2 with probability at least $1-\gamma$.

Theorem 4.4. *Let* G *be an undirected weighted connected graph and* $\lambda \in (0,2]$ *be a lower bound on* λ_2. *Algorithm* $\mathcal{B}$ *outputs a* δ*-multiplicative approximation of* λ_2 *with probability at least* $1-\gamma$, *and uses* $O(\frac{1}{\delta}\log(n\mathfrak{d}/\lambda) + \log\log(1/\gamma))$ *space.*

Proof. Let us first analyze Algorithm $\mathcal{B}$ under the assumption that at Steps 5, 6 and 12, Procedure ESTIMATE-NORM(G,k,v,ϵ,ζ) always correctly outputs an ϵ-additive approximation of $\|M_G^k v\|$. Then during the execution of the algorithm, C_1 and C_2 are ϵ-additive approximations of $\|M_G^k v\|$ and $\|M_G^{k+1} v\|$, respectively.

Lemma 4.2 guarantees that for each $v \in \Sigma$, the inequality $\|M_G^k v\| < \tau$ holds for all integers $k > 2\log(1/\tau)/\lambda_2$, in which case we have $C_1 < \tau + \epsilon < 3\tau/2$ (the same inequality holds for C_2) since $\epsilon < \tau/2$. For each $v \in \Sigma$, the loop of Steps 7-12 is thus repeated at most $2\log(1/\tau)/\lambda_2$ times.

Whenever the ratio C_2/C_1 is computed at Step 8-9, we have $C_1 > 3\tau/2$ and $C_2 > 3\tau/2$. Since $\epsilon < \tau/2$, this means that $\|M_G^k v\| > \tau$ and $\|M_G^{k+1} v\| > \tau$. In this case the quantities C_1 and C_2 are

thus also $\frac{\epsilon}{\tau}$-multiplicative approximations of $\|M_G^k v\|$ and $\|M_G^{k+1} v\|$, respectively. From part (i) of Proposition 4.1 with $\zeta = \epsilon/\tau$, we conclude that $2(1 - R_{\max}) \geq (1 - \delta)\lambda_2$.

Let v be a good vector. Observe that with the choice of τ made at Step 1 we have

$$\frac{\log(1/\tau) - \log(2\sqrt{2}n\mathfrak{d})}{2\lambda_2} = \frac{4\log(\sqrt{2}n\mathfrak{d})}{\delta\lambda_2} \geq \frac{3\log(\sqrt{2}n\mathfrak{d})}{\delta\lambda_2} + 1$$

for $n \geq 4$. There thus necessarily exists at least one integer k satisfying both the condition of Part (ii) of Proposition 4.1 and Inequality (8). From Lemma 4.2, for such a k we have $\|M_G^{k+1} v\| \geq 2\tau$, which gives the lower bound $C_2 \geq 2\tau - \epsilon \geq 3\tau/2$ (the same inequality holds for C_1), and implies that the ratio C_2/C_1 is computed at Steps 8-9. Part (ii) of Proposition 4.1 and Lemma 4.1 thus imply $2(1 - R_{\max}) \leq (1 + \delta)\lambda_2$, as claimed.

From Theorem 4.3, each call of Procedure Estimate-norm requires at most

$$O(\log(nk\mathfrak{d}/\epsilon) + \log\log(1/\zeta)) = O(\frac{1}{\delta}\log(n\mathfrak{d}/\lambda) + \log\log(1/\gamma))$$

space. This bound is also an upper bound on the space complexity of all the other computational steps of Algorithm $\mathcal{B}$. Let us conclude by discussing the success probability of this algorithm. Each application of Estimate-norm errs with probability at most ζ, from Theorem 4.3. There are at most $2 + 2\log(1/\tau)/\lambda_2$ calls to this procedure. The success probability is thus at least $1 - |\Sigma|\,(2 + 2\log(1/\tau)/\lambda)\,\zeta \geq 1 - \gamma$. □

Theorem 4.4 implies Theorem 1.2 by observing again that $\mathfrak{d}/\lambda$ can be upper bounded by a polynomial in n when the weights are polynomially bounded (if G is not connected then we apply Theorem 4.4 on each connected component and taking the minimum of the estimations obtained).

Acknowkedgments

The author is grateful to Richard Cleve, Hirotada Kobayashi, Harumichi Nishimura, Suguru Tamaki and Ryan Williams for helpful comments. This work is supported by the Grant-in-Aid for Young Scientists (A) No. 16H05853, the Grant-in-Aid for Scientific Research (A) No. 16H01705, and the Grant-in-Aid for Scientific Research on Innovative Areas No. 24106009 of the Japan Society for the Promotion of Science and the Ministry of Education, Culture, Sports, Science and Technology in Japan.

References

[1] Stuart J. Berkovitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, pages 147–150, 1984.

[2] Allan Borodin, Joachim von zur Gathen, and John E. Hopcroft. Fast parallel matrix and GCD computations. *Information and Control*, 52(3):241–256, 1982.

[3] Fan R. K. Chung. *Spectral Graph Theory*. American Mathematical Society, 1997.

[4] Fan R. K. Chung and Olivia Simpson. Solving local linear systems with boundary conditions using heat kernel Pagerank. *Internet Mathematics*, 11:4–5, 2015.

[5] Laszlo Csanky. Fast parallel matrix inversion algorithms. In *Proceedings of the 16th Annual Symposium on Foundations of Computer Science*, pages 11–12, 1975.

[6] Dean Doron and Amnon Ta-Shma. On the problem of approximating the eigenvalues of undirected graphs in probabilistic logspace. In *Proceedings of the 42nd International Colloquium*, pages 419–431, 2015.

[7] Peter W. Glynn. Upper bounds on Poisson tail probabilities. *Operation Research Letters*, 6(1):9–14, 1987.

[8] Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. In *Proceedings of the 45th Symposium on Theory of Computing*, pages 911–920, 2013.

[9] Dragoslav Mitrinović. *Analytic Inequalities*. Springer, 1970.

[10] Michael Mitzenmacher and Eli Upfal. *Probability and Computing*. Cambridge University Press, 2005.

[11] Richard Peng and Daniel A. Spielman. An efficient parallel solver for SDD linear systems. In *Proceedings of the 46th Symposium on Theory of Computing*, pages 333–342, 2014.

[12] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual Symposium on Theory of Computing*, pages 81–90, 2004.

[13] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.

[14] Daniel A. Spielman and Shang-Hua Teng. A local clustering algorithm for massive graphs and its application to nearly linear time graph partitioning. *SIAM Journal on Computing*, 42(1):1–26, 2013.

[15] Daniel A. Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.

[16] Amnon Ta-Shma. Inverting well conditioned matrices in quantum logspace. In *Proceedings of the 45th Symposium on Theory of Computing*, pages 881–890, 2013.

[17] Nisheeth K. Vishnoi. *$Lx = b$ — Laplacian Solvers and their Algorithmic Applications*. Now publishers, 2013.

Appendix: Estimating $\left\|M_G^k v\right\|$

In this appendix we explain how to space-efficiently estimate $\left\|M_G^k v\right\|$ for any vector $v \in \Sigma$, and prove Theorem 4.3.

The following lemma first shows how to space-efficiently estimate the quantity $2^{-k}\binom{k}{s}$.

Lemma 4.5. *Let δ and ζ be any real numbers such that $\delta \in (0,1)$ and $\zeta > 0$. There exists a $O(\log(ks/\delta) + \log\log(1/\zeta))$-space algorithm that, when given as input two integers $k \geq 1$ and $s \in \{0,\dots k\}$, outputs an δ-additive approximation of $2^{-k}\binom{k}{s}$ with probability $1-\zeta$.*

Proof. Our algorithm is as follows. The algorithm creates a counter C initialized to zero, and repeat $m = \left\lceil \frac{\log(2/\zeta)}{2\delta^2} \right\rceil$ times the following: take s bits uniformly at random and increment C by one if exactly k among these s bits are one. The algorithm finally outputs C/m. Observe that this algorithm can be implemented in space linear in $\log(1/\delta)$, $\log k$, $\log s$ and $\log\log(1/\zeta)$.

The expected value of the output of this algorithm is precisely $2^{-k}\binom{k}{s}$. From Chernoff bound, the output of the algorithm is thus a δ-additive approximation of this quantity with probability at least $1 - 2e^{-2m\delta^2} \geq 1-\zeta$. □

The procedure ESTIMATE-NORM(G,k,v,ϵ,γ) estimating $\left\|M_G^k v\right\|$ is described in Figure 3.

We now prove Theorem 4.3 stated in Section 4.

Proof of Theorem 4.3. Observe that

$$M_G^k = \sum_{i=1}^{n} (1-\lambda_i/2)^k u_i^t u_i = D_G^{1/2}\left(\sum_{s=0}^{k} \frac{\binom{k}{s}}{2^k} P_G^s\right) D_G^{-1/2}.$$

Moreover M_G^k is symmetric. Let us write $v = \sqrt{\frac{d_{\ell_2}}{d_{\ell_1}+d_{\ell_2}}}e_{\ell_1} - \sqrt{\frac{d_{\ell_1}}{d_{\ell_1}+d_{\ell_2}}}e_{\ell_2}$ as in Step 2 of the procedure. We thus have

$$\left\|M_G^k v\right\| = \left\| v^t D_G^{1/2}\left(\sum_{s=0}^{k} \frac{\binom{k}{s}}{2^k} P_G^s\right) D_G^{-1/2}\right\| = \sqrt{\sum_{i=1}^{n} \Gamma_i^2}$$

where

$$\Gamma_i = \sum_{s=0}^{k} \frac{\binom{k}{s}}{2^k}\left(\sqrt{\frac{d_{\ell_1}d_{\ell_2}}{(d_{\ell_1}+d_{\ell_2})d_i}}\langle e_i, e_{\ell_1}P_G^s\rangle - \sqrt{\frac{d_{\ell_1}d_{\ell_2}}{(d_{\ell_1}+d_{\ell_2})d_i}}\langle e_i, e_{\ell_2}P_G^s\rangle\right).$$

Note that $|\Gamma_i| \leq 1$ for any $i \in \{1,\dots,n\}$.

At the end of Steps 9-11 we have

$$\Pr\left[|S/r - \langle e_i, e_{\ell_1}P_G^s\rangle| \leq \delta\right] \geq 1 - 2e^{-2r\delta^2} \geq 1-\zeta, \tag{9}$$

using the same argument as in the analysis in Theorem 3.3. The same bound, with ℓ_1 replaced by ℓ_2, holds at the end of Steps 14-16 as well. Lemma 4.5 also shows that a is a δ-additive approximation of $2^{-k}\binom{k}{s}$ with probability at least $1-\zeta$ at Step 7. Let us continue our analysis

Input: an undirected, weighted and connected graph G on n vertices,
a vector $v \in \Sigma$, a positive integer k,
a precision parameter $\epsilon \in (0,1]$, an error parameter $\gamma \in (0,1]$

1 $\delta \leftarrow \frac{\epsilon^2\sqrt{2}}{54(k+1)n\mathfrak{d}}$; $\zeta \leftarrow \frac{\gamma}{3n(k+1)}$; $r \leftarrow \left\lceil \frac{\log(2/\zeta)}{2\delta^2} \right\rceil$;
2 Let (ℓ_1, ℓ_2) be the indices such that $v = \sqrt{\frac{d_{\ell_2}}{d_{\ell_1}+d_{\ell_2}}} e_{\ell_1} - \sqrt{\frac{d_{\ell_1}}{d_{\ell_1}+d_{\ell_2}}} e_{\ell_2}$;
3 $R \leftarrow 0$;
4 **for** i from 1 to n
5 $\quad Q \leftarrow 0$;
6 $\quad$ **for** s from 0 to k
7 $\quad\quad$ Compute an estimate a of $2^{-k}\binom{k}{s}$ using Lemma 4.5 with δ and ζ;
8 $\quad\quad S \leftarrow 0$;
9 $\quad\quad$ **repeat** r times:
10 $\quad\quad\quad$ Run the walk P_G starting on vertex ℓ_1 for s steps;
11 $\quad\quad\quad$ If the walk ends on vertex i then $S \leftarrow S + 1$;
12 $\quad\quad Q \leftarrow Q + \frac{aS}{r}\sqrt{\frac{d_{\ell_1} d_{\ell_2}}{(d_{\ell_1}+d_{\ell_2})d_i}}$;
13 $\quad\quad S \leftarrow 0$;
14 $\quad\quad$ **repeat** r times:
15 $\quad\quad\quad$ Run the walk P_G starting on vertex ℓ_2 for s steps;
16 $\quad\quad\quad$ If the walk ends on vertex i then $S \leftarrow S + 1$;
17 $\quad\quad Q \leftarrow Q - \frac{aS}{r}\sqrt{\frac{d_{\ell_1} d_{\ell_2}}{(d_{\ell_1}+d_{\ell_2})d_i}}$;
18 $\quad R \leftarrow R + Q^2$;
19 Output $\sqrt{R}$.

Figure 3: Procedure ESTIMATE-NORM$(G, k, v, \epsilon, \gamma)$ computing an ϵ-additive approximation of $\left\|M_G^k v\right\|$ with probability at least $1 - \gamma$.

under the assumption that all these approximations are correct (we discuss the overall success probability at the end of the proof). At Step 12 we thus have

$$\left| \frac{aS}{r} - \frac{\binom{k}{s}}{2^k} \langle e_i, e_{\ell_1} P_G^s \rangle \right| \leq \left(\delta + \frac{\binom{k}{s}}{2^k} + \langle e_i, e_{\ell_1} P_G^s \rangle \right) \delta \leq 3\delta,$$

and the same bound, with ℓ_1 replaced by ℓ_2, holds at Step 17 as well. For any $i \in \{1, \ldots, n\}$ we thus have

$$|Q - \Gamma_i| \leq 6(k+1)\delta \max_{\ell_1, \ell_2 \in \{1,\ldots,n\}} \sqrt{\frac{d_{\ell_1} d_{\ell_2}}{(d_{\ell_1}+d_{\ell_2})d_i}} \leq \frac{6}{\sqrt{2}}(k+1)\delta\mathfrak{d}.$$

at the end of the loop of Steps 5-17, which implies

$$\left|Q^2 - \Gamma_i^2\right| = |Q - \Gamma_i| \times |Q + \Gamma_i| \leq \frac{18}{\sqrt{2}}(k+1)\delta\mathfrak{d}.$$

since $|\Gamma_i| \leq 1$ and $|Q| \leq |Q - \Gamma_i| + |\Gamma_i| \leq 2$. We thus have

$$\left| R - \left\| M_G^k v \right\|^2 \right| \leq \frac{18}{\sqrt{2}}(k+1)n\delta\mathfrak{d} \leq \epsilon^2/3.$$

at the end of of the algorithm. Let us show that $\sqrt{R}$ is an ϵ-additive estimation of $\left\| M_G^k v \right\|$ by considering two cases. In the case $\sqrt{|R|} \leq \epsilon/3$ we get

$$\left| \sqrt{R} - \left\| M_G^k v \right\| \right| \leq \sqrt{R} + \left\| M_G^k v \right\| \leq \sqrt{R} + \sqrt{R + \left| R - \left\| M_G^k v \right\|^2 \right|} \leq \epsilon/3 + \sqrt{\epsilon^2/9 + \epsilon^2/3} = \epsilon.$$

Now in the case $\sqrt{|R|} \geq \epsilon/3$ we get

$$\left| \sqrt{R} - \left\| M_G^k v \right\| \right| = \frac{\left| R - \left\| M_G^k v \right\|^2 \right|}{\left| \sqrt{R} + \left\| M_G^k v \right\| \right|} \leq \frac{\epsilon^2/3}{\epsilon/3} = \epsilon.$$

The space complexity of this algorithm is $O(\log(nk/\epsilon) + \log\log\gamma)$, from Lemma 4.5 and the observation that all other computational steps can be implemented with registers of this size. Let us conclude by discussing the success probability of this algorithm. Errors can only occur at Steps 6 or 10. From Inequality (9) and Lemma 4.5, and using the union bound, we know that the overall success probability is at least $1 - 3\zeta n(k+1) = 1 - \gamma$. □